An Approach to Qualitative Belief Change Modulo Ontic Strength

Gavin Rens[1] and **Gabriele Kern-Isberner**[2]

Abstract. Sometimes, strictly choosing between belief revision and belief update is inadequate in a dynamical, uncertain environment. Boutilier combined the two notions to allow updates in response to external changes to inform an agent about its prior beliefs. His approach is based on ranking functions. Rens proposed a new method to trade off probabilistic revision and update, in proportion to the agent's confidence for whether to revise or update. In this paper, we translate Rens's approach from a probabilistic setting to a setting with ranking functions. Given the translation, we are able to compare Boutilier's and Rens's approaches. We found that Rens's approach is an extension of Boutilier's.

1 Introduction

Traditionally, *belief revision* is regarded as change of beliefs about the objective static state of the world. And *belief update* is regarded as change of beliefs due to recording the change which occurred in the underlying state of the world. We shall use the generic term *belief change* to including belief update and belief revision. An agent may not always be certain whether an observation is a side-effect of an action/event (requiring update), or whether the observation did not have a physical cause and is thus pure information (requiring revision).

Boutilier [3] proposed a generalized qualitative update procedure, which combines both belief update and revision. He used ranking functions as advocated by Spohn [13, 14] to capture notions of preference. Rens [11] proposed a quantitative approach to mix probabilistic belief update and revision, where the trade-off is controlled by the so-called *ontic strength* of the observation received. To our knowledge, his "mixture" method is novel. In this paper, we propose a translation of Rens's method back to a qualitative setting using Spohn-rankings. The difference in the present approach to that of Boutilier is that ours trades belief update and revision off in proportion to the agent's judgement of the ontic strength of the received evidence.

There are several reason why we would like a qualitative version of Rens's hybrid stochastic belief change (HSBC).

- Ordering preferred worlds by ranking them instead of providing exact probabilities may be more intuitive for agent designers.
- Some domains may not require the agent to work with precise values like probabilities, and computations over ranked preferences are then cheaper, because finding the minimum of a set is generally cheaper than finding its sum (the distinction between minimization and summation will become clear later).
- We may gain insights about the relationship between belief revision and belief update when analysed in the qualitative belief change setting.

Let L be the classical propositional language, and W the (finite) set of possible worlds (valuations) induced from a finite set of propositional variables. We denote the models of a sentence $\alpha \in L$ by $[\![\alpha]\!]$ and the fact that w satisfied α by $w \Vdash \alpha$. For a set of sentences $K \subseteq L$, $[\![K]\!] := \{w \in W \mid \forall \beta \in K, w \Vdash \beta\}$. We refer to a probability function or a ranking function as an *epistemic state*. In this paper, we denote the result of a belief change operation as $\epsilon \circ \alpha$, where ϵ is an epistemic state and $\circ$ is the operator. If we need to refer to the value of a particular world w in the changed epistemic state, we write $\epsilon^{\circ}_{\alpha}(w)$.

Next, we review the essentials of Rens's HSBC construction. In Section 3, we provide the qualitative, rank-based translation of HSBC (i.e., HQBC). We analyse our HQBC with respect to two fundamental classical rationality postulates in Section 4. In Section 5, we compare our hybrid qualitative belief change construction to Boutilier's *generalized update* construction. Two examples are presented in Section 6, and we end the paper with a summary of what has been achieved here, and a discussion about related and future work.

2 Hybrid Belief Change via Probability Theory

Rens [11] proposed the *hybrid stochastic belief change* (HSBC) operation to combine notions of probabilistic belief revision and probabilistic belief update. HSBC may be employed in agents who deal with uncertainty by maintaining a probability distribution b over possible worlds w they could be in. That is, $b : W \to [0, 1]$, such that $\sum_{w \in W} b(w) = 1$, and $b(\alpha) := \sum_{w \in W, w \Vdash \alpha} b(w)$ for all $\alpha \in L$. b will often be represented as a set of pairs $\{(w, p) \mid w \in W, p \in [0, 1]\}$. We refer to b as an epistemic state in the context of this work. In the HSBC framework, an agent maintains an epistemic state, which changes as new information is received or observed.

Rens [11] proposes the tuple $\langle W, Evt, T, E, O, os \rangle$ to formalize the HSBC framework, where

- W is a set of possible worlds;
- Evt is a set of atomic events;
- $T : W \times Evt \times W \to [0, 1]$ is a *transition function* such that for every $e \in Evt$ and $w \in W$, $\sum_{w' \in W} T(w, e, w') = 1$, and $T(w, e, w')$ models the probability of a transition to world w', given the occurrence of event e in world w;
- E is the *event function* such that $E(e, w) = P(e \mid w)$, the probability of the occurrence of event e in w;

[1] Centre for AI Research, and University of KwaZulu-Natal, School of Mathematics, Statistics and Computer Science, and CSIR Meraka, South Africa
[2] Dortmund University of Technology, Dortmund, Germany

- $O : L \times W \to [0,1]$ is an *observation function* such that for every world w, $\sum_{\alpha \in \Omega} O(\alpha, w) = 1$, and $O(\alpha, w)$ models the probability of observing α in w, where $\Omega \subset L$ is the set of possible observations, up to equivalence, and where if $\alpha \equiv \beta$, then $O(\alpha, w) = O(\beta, w)$, for all worlds w;[3]
- $os : \Omega \times W \to [0,1]$ such that $os(\alpha, w)$ is the agent's *ontic strength* for α perceived in w.

In HSBC, the epistemic state updated with α (denoted $b \diamond \alpha$) is defined as

$$b \diamond \alpha := \big\{(w', p') \mid w' \in W, p' = \frac{1}{\gamma} O(\alpha, w') \sum_{w \in W} \sum_{e \in Evt} T(w, e, w') E(e, w) b(w)\big\},$$

where γ is a normalizing factor.

It is mostly agreed upon that Bayesian conditioning corresponds to classical belief *expansion*. This is evidenced by Bayesian conditioning (BC) being defined only when $b(\alpha) \neq 0$ (i.e., when α does not contradict the agent's current beliefs). In other words, one could define revision to be

$$b \,\mathsf{BC}\, \alpha := \{(w, p) \mid w \in W, p = P(w \mid \alpha)\},$$

as long as $P(\alpha) \neq 0$, where

$$P(w \mid \alpha) := \frac{O(\alpha, w) b(w)}{\sum_{w' \in W} O(\alpha, w') b(w')}.^{4} \qquad (1)$$

To accommodate cases where $b(\alpha) = 0$, that is, where α contradicts the agent's current beliefs and its beliefs need to be revised in the stronger sense, we shall make use of *imaging*. Imaging was introduced by Lewis [9] as a means of revising a probability function. Informally, Lewis's original solution for accommodating contradicting evidence α is to move the probability of each world to its closest, α-world. Lewis made the strong assumption that every world has a *unique* closest α-world. More general versions of imaging allow worlds to have *several*, equally proximate closest worlds.

In two papers, Rens [11] and colleagues [12] propose *generalized imaging*: Let $Min(\alpha, w, d)$ be the set of α-worlds closest to w measured with d, some acceptable measure of distance between worlds (e.g., Hamming or Dalal distance). Formally,

$$Min(\alpha, w, d) := \{w' \in [\![\alpha]\!] \mid \forall w'' \in [\![\alpha]\!], d(w, w') \leq d(w, w'')\}.$$

Then generalized imaging (denoted GI) is defined as

$$b \,\mathsf{GI}\, \alpha := \big\{(w, p) \mid w \in W, p = 0 \text{ if } w \notin [\![\alpha]\!], \text{ else } p = \sum_{\substack{w' \in W \\ w \in Min(\alpha, w', d)}} b(w')/|Min(\alpha, w', d)|\big\}.$$

Rens [11] argues that if observation likelihoods are known, they should be used to weight the probabilities computed by the GI operation; a new imaging operation is thus defined as

$$b \,\mathsf{OGI}\, \alpha := \Big\{(w, p) \mid w \in W, p = \frac{O(\alpha, w) b^{\mathsf{GI}}_{\alpha}(w)}{\sum_{w' \in W} O(\alpha, w') b^{\mathsf{GI}}_{\alpha}(w')}\Big\},$$

where the denominator is a normalizing factor. At last, with respect to revision, Rens [11] defines BCI to revise by conditioning when the evidence does not contradict the agent's current beliefs, and to revise by imaging otherwise:

$$b \,\mathsf{BCI}\, \alpha := \begin{cases} b \,\mathsf{BC}\, \alpha & \text{if } b(\alpha) > 0 \\ b \,\mathsf{OGI}\, \alpha & \text{if } b(\alpha) = 0 \end{cases}$$

Finally, Rens [11] proposes a way of trading off the probabilistic update and probabilistic revision, using the notion of ontic strength. He argues that an agent could reason with a range of degrees for information being ontic (the effect of a physical action or occurrence) or epistemic (purely informative). It is assumed that the higher the information's degree of being ontic, the lower the epistemic status of that information. "An agent has a certain sense of the degree to which a piece of received information is due to a physical action or event in the world. This sense may come about due to a combination of sensor readings and reasoning. If the agent performs an action and a change in the local environment matches the expected effect of the action, it can be quite certain that the effect is ontic information," [11, p. 129].

$os(\alpha, w)$ is defined to equal 1 when α is certainly ontic in w, and 0 when α is certainly epistemic (the epistemic strength of α in w is $es(\alpha, w) := 1 - os(\alpha, w)$).

The hybrid stochastic change of epistemic state b due to new information α with ontic strength (denoted $b \lhd \alpha$) is defined as

$$b \lhd \alpha := \Big\{(w, p) \mid w \in W, p = \frac{1}{\gamma'}\big[es(\alpha, w) b^{\mathsf{BCI}}_{\alpha}(w) + os(\alpha, w) b^{\diamond}_{\alpha}(w)\big]\Big\},$$

where γ' is a normalizing factor so that $\sum_{w \in W} b^{\lhd}_{\alpha}(w) = 1$.

3 Hybrid Belief Change via Ranking Theory

Let κ be a ranking on worlds in W, representing the agent's current epistemic state, as first proposed by Spohn [13]. That is, $\kappa : W \to \mathbb{N} \cup \{\infty\}$, where $\mathbb{N} = \{0, 1, 2, \ldots\}$, such that there exists a $w \in W$ for which $\kappa(w) = 0$ and $\kappa(w_i) \leq \kappa(w_j)$ is interpreted as world w_i being at least as plausible or preferred as world w_j. $\kappa(w^{\times}) = \infty$ is meant to indicate that $w^{\times}$ is impossible, implausible, least preferred. Worlds w' for which $\kappa(w') = 0$ are considered most plausible, most preferred, or believed. In fact, ranking functions are rankings of *im*plausibility. The degree of plausibility of proposition α is

$$\kappa(\alpha) := \min_{w \in W, w \Vdash \alpha} \{\kappa(w)\}. \qquad (2)$$

We shall denote an agent's belief set, given epistemic state κ, as

$$Bel(\kappa) := \{\beta \in L \mid \kappa^{-1}(0) \subseteq [\![\beta]\!]\}.$$

Since Spohn's ranking functions can be considered as the logarithm of probabilities [7], when translating from probability theory to Spohn ranking theory, multiplication becomes addition, division becomes subtraction, and summation ($\sum_{w \in W}$) becomes minimization ($\min_{w \in W}$).

Conditional plausibility is defined as [13]

$$\kappa(\beta \mid \alpha) := \kappa(\alpha \wedge \beta) - \kappa(\alpha).$$

Let ϕ_w be a complete theory for w. It will be useful to know that $\kappa(\alpha \wedge \beta) = \kappa(\alpha \mid \beta) + \kappa(\beta)$. And consequently, that $\kappa(\alpha) = \min_{w \in W}\{\kappa(\alpha \mid \phi_w) + \kappa(w)\}$. One can also define $\kappa(w \mid \alpha)$ in terms of $\kappa(\alpha \mid w)$ as follows.

$$\begin{aligned} \kappa(w \mid \alpha) &= \kappa(\phi_w \wedge \alpha) - \kappa(\alpha) \\ &= \kappa(\alpha \mid \phi_w) + \kappa(w) - \kappa(\alpha). \end{aligned}$$

---

[3] $\equiv$ denotes logical equivalence.

[4] Note that $b(\alpha)$ is equivalent to $P(\alpha)$.

A direct translation of the Bayes Rule would suggest that the above result is analogous to that rule.

Definition 1. *The tuple $\langle W, Evt, T_Q, E_Q, O_Q, os\rangle$ is a hybrid qualitative belief change (HQBC) model, where*

- *W is a set of possible worlds;*
- *Evt is a set of atomic events;*
- *$T_Q : W \times Evt \times W \to \mathbb{N} \cup \{\infty\}$ is a* transition ranking *such that for every $w \in W$ and $e \in Evt$, $\min_{w' \in W}\{T_Q(w, e, w')\} = 0$ and $T_Q(w, e, w')$ models the plausibility of a transition to world w', given the occurrence of event e in world w;*
- *$E_Q : Evt \times W \to \mathbb{N} \cup \{\infty\}$ is the* event ranking *such that for every $w \in W$, $\min_{e \in Evt}\{E_Q(e, w)\} = 0$ and $E_Q(e, w)$ models the plausibility of the occurrence of event e in w;*
- *$O_Q : L \times W \to \mathbb{N} \cup \{\infty\}$ is an* observation ranking *such that for every world w, $\min_{\alpha \in L}\{O(\alpha, w)\} = 0$ and $O_Q(\alpha, w)$ models the plausibility of observing α in w and where if $\alpha \equiv \beta$, then $O_Q(\alpha, w) = O_Q(\beta, w)$, for all worlds w*[5];[6]
- *$os : L \times W \to \mathbb{N}$, where $os(\alpha, w)$ is the agent's* ontic strength *for α perceived in w (note that os is not a κ-function).*

In the qualitative version of the hybrid belief change framework, epistemic strength is defined as the complement of ontic strength. Unfortunately, the notion of complement is not strictly defined for ranking theory. We thus define epistemic strength as the complement of ontic strength with respect to a 'top' value.

Definition 2. *Let τ be an even number in $\mathbb{N}$, but do not let $\tau = \infty$. Epistemic strength is defined as the τ-complement of os.*

$$es(\alpha, w) := \tau - os(\alpha, w).$$

for all possible observations α and for all worlds $w \in W$.

To specify that the agent has no preference for an observation being ontic or epistemic, choose $os(\alpha, w) = \tau/2$ for all w. Then $es(\alpha, w) = \tau/2$ for all w.[7]

Let κ be regarded as an agent's epistemic state and α a new piece of information to be accommodated. We can define the operation which revises an epistemic state using conditional plausibility:

$$\kappa\ CP\ \alpha =: \{(w, n) \mid w \in W, n = Q(w \mid \alpha)\},$$

as long as $\kappa(\alpha) \neq \infty$, where

$$Q(w \mid \alpha) := O_Q(\alpha, w) + \kappa(w) - \min_{w' \in W}\{O_Q(\alpha, w') + \kappa(w')\}$$

is justified by the translation of $P(w \mid \alpha)$ (Eq. 1) from probability theory to ranking theory. The definition of $Q(w \mid \alpha)$ can also be derived from first principles, which we leave out here.

As with probabilistic conditionalization, plausibilistic conditionalization is undefined when the evidence/observation is inconsistent with the agent's current epistemic state. A plausibilistic version of imaging can deal with this problem in the qualitative setting:

Translate b GI α to

$$\kappa\ \mathsf{GI}\ \alpha := \Big\{(w, n) \mid w \in W, n = \infty \text{ if } w \notin [\![\alpha]\!], \text{ else } n = \min_{\substack{w' \in W \\ w \in Min(\alpha, w', d)}}\{\kappa(w')\}\Big\}.$$

[5] $\equiv$ denotes logical equivalence.
[6] $\equiv$ denotes logical equivalence.
[7] Due to τ being even, $\tau/2$ is guaranteed to be a whole number.

Example 1. *Let the vocabulary be $\{q, r, s\}$ and the current epistemic state $\kappa_1 = \{(\bar{q}\bar{r}\bar{s}, 0), (\bar{q}r\bar{s}, 1), (q\bar{r}\bar{s}, 2), (qr\bar{s}, 3), (qrs, \infty), (q\bar{r}s, \infty), (\bar{q}rs, \infty), (\bar{q}\bar{r}s, \infty)\}$. Let d be defined as Hamming distance. Suppose the observation α received is $(q \wedge r) \vee (q \wedge \neg r \wedge s)$. Then*

$$\begin{array}{ll} Min(\alpha, qrs, d) = \{qrs\} & Min(\alpha, qr\bar{s}, d) = \{qr\bar{s}\} \\ Min(\alpha, q\bar{r}s, d) = \{q\bar{r}s\} & Min(\alpha, q\bar{r}\bar{s}, d) = \{qr\bar{s}, q\bar{r}s\} \\ Min(\alpha, \bar{q}rs, d) = \{qrs\} & Min(\alpha, \bar{q}r\bar{s}, d) = \{qr\bar{s}\} \\ Min(\alpha, \bar{q}\bar{r}s, d) = \{q\bar{r}s\} & Min(\alpha, \bar{q}\bar{r}\bar{s}, d) = \{qr\bar{s}, q\bar{r}s\} \end{array}$$

and

- $\kappa^{\mathsf{GI}}_{1\alpha}(qrs) = \min\{\kappa(qrs), \kappa(\bar{q}rs)\} = \min\{\infty, \infty\} = \infty$,
- $\kappa^{\mathsf{GI}}_{1\alpha}(qr\bar{s}) = \min\{\kappa(qr\bar{s}), \kappa(q\bar{r}\bar{s}), \kappa(\bar{q}r\bar{s}), \kappa(\bar{q}\bar{r}\bar{s})\}$ $= \min\{3, 2, 1, 0\} = 0$,
- $\kappa^{\mathsf{GI}}_{1\alpha}(q\bar{r}s) = \min\{\kappa(q\bar{r}s), \kappa(q\bar{r}\bar{s}), \kappa(\bar{q}\bar{r}s), \kappa(\bar{q}\bar{r}\bar{s})\}$ $= \min\{\infty, 2, \infty, 0\} = 0$,
- $\kappa^{\mathsf{GI}}_{1\alpha}(q\bar{r}\bar{s}) = \kappa^{\mathsf{GI}}_{1\alpha}(\bar{q}rs) = \kappa^{\mathsf{GI}}_{1\alpha}(\bar{q}r\bar{s}) = \kappa^{\mathsf{GI}}_{1\alpha}(\bar{q}\bar{r}s) = \kappa^{\mathsf{GI}}_{1\alpha}(\bar{q}\bar{r}\bar{s})$ $= \infty$.

Notice that $(q \wedge r) \vee (q \wedge \neg r \wedge s)$ does not contradict κ_1 (i.e., $\kappa_1((q \wedge r) \vee (q \wedge \neg r \wedge s)) \neq \infty$). To show that qualitative imaging can deal with observations contradicting the agent's epistemic state, consider the following example.

Example 2. *We consider the same setting as in Example 1. Suppose the observation β received is $q \wedge s$. Note that $\kappa_1(\beta) = \infty$, that is, $q \wedge s$ is deemed impossible in κ_1. Then*

- $\kappa^{\mathsf{GI}}_{1\beta}(qrs) = \min\{\kappa(qrs), \kappa(qr\bar{s}), \kappa(\bar{q}rs), \kappa(\bar{q}r\bar{s})\} = \min\{\infty, 3, \infty, 1\} = 1$,
- $\kappa^{\mathsf{GI}}_{1\beta}(q\bar{r}s) = \min\{\kappa(q\bar{r}s), \kappa(q\bar{r}\bar{s}), \kappa(\bar{q}\bar{r}s), \kappa(\bar{q}\bar{r}\bar{s})\}$ $= \min\{\infty, 2, \infty, 0\} = 0$,
- $\kappa^{\mathsf{GI}}_{1\beta}(qr\bar{s}) = \kappa^{\mathsf{GI}}_{1\beta}(q\bar{r}\bar{s}) = \kappa^{\mathsf{GI}}_{1\beta}(\bar{q}rs) = \kappa^{\mathsf{GI}}_{1\beta}(\bar{q}r\bar{s}) = \kappa^{\mathsf{GI}}_{1\beta}(\bar{q}\bar{r}s) = \kappa^{\mathsf{GI}}_{1\beta}(\bar{q}\bar{r}\bar{s}) = \infty$.

Now, qualitative generalized imaging can be weighted/modulated by the plausibility of the evidence in a particular world:

$$\kappa\ \mathsf{OGI}\ \alpha := \{(w, n) \mid w \in W, n = \kappa^{\mathsf{GI}}_{\alpha}(w) + O_Q(\alpha, w) - \delta\},$$

where δ is a normalization factor defined as

$$\delta := \min_{w \in W}\{\kappa^{\mathsf{GI}}_{\alpha}(w) + O_Q(\alpha, w)\}.$$

Example 3. *Continuing with the previous examples, suppose $O_Q(q \wedge s, \bar{q}\bar{r}\bar{s}) = 1$ and for all $w \neq \bar{q}\bar{r}\bar{s}$, $O_Q(q \wedge s, w) = 0$. Then*

- $\kappa^{\mathsf{OGI}}_{1\beta}(qrs) = \min\{\kappa(qrs)+0, \kappa(qr\bar{s})+0, \kappa(\bar{q}rs)+0, \kappa(\bar{q}r\bar{s})+0\} - \delta = \min\{\infty+0, 3+0, \infty+0, 1+0\} - \delta = 1 - 1 = 0$,
- $\kappa^{\mathsf{OGI}}_{1\beta}(q\bar{r}s) = \min\{\kappa(q\bar{r}s)+0, \kappa(q\bar{r}\bar{s})+0, \kappa(\bar{q}\bar{r}s)+0, \kappa(\bar{q}\bar{r}\bar{s})+1\} - \delta = \min\{\infty+0, 2+0, \infty+0, 0+1\} - \delta = 1 - 1 = 0$,
- $\kappa^{\mathsf{OGI}}_{1\beta}(qr\bar{s}) = \kappa^{\mathsf{OGI}}_{1\beta}(q\bar{r}\bar{s}) = \kappa^{\mathsf{OGI}}_{1\beta}(\bar{q}rs) = \kappa^{\mathsf{OGI}}_{1\beta}(\bar{q}r\bar{s}) = \kappa^{\mathsf{OGI}}_{1\beta}(\bar{q}\bar{r}s) = \kappa^{\mathsf{OGI}}_{1\beta}(\bar{q}\bar{r}\bar{s}) = \infty$.

In Example 2, $q\bar{r}s$ is most plausible in $\kappa^{\mathsf{GI}}_{1\beta}$, but in Example 3, due to $q \wedge s$ being slightly less plausibly perceived in $\bar{q}\bar{r}\bar{s}$ than in any other world, $\bar{q}\bar{r}\bar{s}$ becomes slightly less plausible in $\kappa^{\mathsf{OGI}}_{1\beta}$. Thus, in

Example 3, qrs and $q\bar{r}s$ share the status of being most plausible in the agent's revised epistemic state.

Finally, a qualitative version of BCI can be defined, which revises by conditional plausibility when the evidence does not contradict the agent's current beliefs, and revises by qualitative imaging otherwise:

$$\kappa \,\mathsf{CPI}\, \alpha := \begin{cases} \kappa \,\mathsf{CP}\, \alpha & \text{if } \kappa(\alpha) \neq \infty \\ \kappa \,\mathsf{OGI}\, \alpha & \text{if } \kappa(\alpha) = \infty \end{cases}$$

Turning now to belief update, $b \diamond \alpha$ is translated to

$$\kappa \diamond \alpha := \Big\{(w', n) \mid w' \in W, n = \quad (3)$$
$$O_Q(\alpha, w') + \min_{\substack{w \in W \\ e \in Evt}} \big\{T_Q(w, e, w') + E_Q(e, w) + \kappa(w)\big\} - \delta'\Big\},$$

where δ' is a normalizing factor defined as

$$\delta' := \min_{w' \in W} \Big\{O_Q(\alpha, w') + \min_{\substack{w \in W \\ e \in Evt}} \big\{T_Q(w, e, w') + E_Q(e, w) + \kappa(w)\big\}\Big\}.$$

Example 4. *We continue, using the vocabulary and epistemic state of the previous examples. For illustrative purposes, we keep the observation, transition and event models very simple, with an arbitrary specification: For all $w \in W$, let $O_Q(\beta, w) = 0$ if $w \Vdash \beta$, else, $O_Q(\beta, w) = 1$. Let there be two events: $Evt = \{e_1, e_2\}$. Let $W = \{w_1, w_2, \ldots, w_n\}$. $E_Q(e_k, w_i) = i \times k$, except for $E_Q(e_1, w_1) = 0$. $T_Q(w_i, e_k, w_j) = i \times j \times k$, except for $T_Q(w_1, e_1, w_1) = 0$. Then the ranks of the first two worlds are*

- $\kappa^{\diamond}_{\beta}(qrs) = O_Q(\beta, qrs) + \min_{\substack{w \in W \\ e \in Evt}} \big\{T_Q(w, e, qrs) + E_Q(e, w) + \kappa(w)\big\} - \delta' = 0 + \min\{0+0+\infty, 2+2+3, \ldots, 16+8+0\} - \delta' = 7 - \delta'$ *and*
- $\kappa^{\diamond}_{\beta}(qr\bar{s}) = O_Q(\beta, qr\bar{s}) + \min_{\substack{w \in W \\ e \in Evt}} \big\{T_Q(w, e, qr\bar{s}) + E_Q(e, w) + \kappa(w)\big\} - \delta' = 1 + \min\{2+1+\infty, 4+2+3, 6+3+\infty, \ldots, 32+16+\infty\} - \delta' = 10 - \delta'$.

We do not work out the ranks of the other worlds.

Given an epistemic state κ and a new observation α, we propose the following HQBC operation.

$$\kappa \lhd \alpha := \big\{(w, n) \mid w \in W, n = \quad (4)$$
$$\min\{\kappa^{CPI}_{\alpha}(w) + es(\alpha, w), \quad \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\} - \delta''\big\},$$

where δ'' is a normalizing factor defined as

$$\delta'' := \min_{w \in W} \big\{ \min\{\kappa^{CPI}_{\alpha}(w) + es(\alpha, w), \quad \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\}\big\}.$$

$\kappa^{\lhd}_{\alpha}(w)$ can be read as 'The rank of w after revision if revision is more plausible given the epistemic strength of α at w, else, the rank of w after update (given the ontic strength).'

4 Analysis of HQBC w.r.t. Rationality Postulates

In this section we shall assess two fundamental postulates generally agreed upon as necessary (but not sufficient) for belief change to be rational [6, e.g.]. The *categorical matching* postulate (CM) states that the representation of an agent's state of knowledge/belief should have the same formal structure before and after the application of the belief change operation under consideration. The *success* postulate (S) states that the observation/evidence with which an agent's state is to be changed should be believed (with certainty) after the belief change operation.[8] In the rest of this section, we assume that $\alpha \in L$ is any logically satisfiable piece of information.

[8] Here it is assumed that the incoming information is certainly correct.

Definition 3. *We say*

- *event e is* possible in κ *iff there exists a world $w \in W$ such that $\kappa(w) \neq \infty$ and $E_Q(e, w) \neq \infty$;*
- *event e is* event-rational *when for all $w \in W$: there exists a w' such that $T_Q(w, e, w') \neq \infty$ iff $E_Q(e, w) \neq \infty$;*
- *evidence α is an e-*signal *when for all $w' \in W$: there exists a w such that $T_Q(w, e, w') \neq \infty$ iff $O_Q(\alpha, w') \neq \infty$;*
- *evidence α is* trustworthy *iff for all $w \in W$, if $w \nVdash \alpha$, then $O_Q(\alpha, w) = \infty$;*
- *evidence α is* clear *iff for all $w \in W$, if $w \Vdash \alpha$, then $O_Q(\alpha, w) = 0$;*
- *evidence α is* weakly observable *iff there exists a $w \in W$ such that $w \Vdash \alpha$ and $O(\alpha, w) \neq \infty$;*
- *evidence α is* strongly observable *iff for all $w \in W$ for which $w \Vdash \alpha$, $O(\alpha, w) \neq \infty$.*

Except for *possibility*, the definitions in the list above are adapted from Rens [11].

Postulate (CM) If κ is a ranking function, then so is $\kappa \lhd \alpha$.

Lemma 1. *If α is strongly observable, then $\kappa \,\mathsf{CPI}\, \alpha$ is a ranking function.*

Proof. Omitted to save space; available on request. □

Lemma 2. *Let the HQBC model be specified such that there exists an event-rational event $e \in Evt$ possible in κ, and α is an e-signal. Then $\kappa \diamond \alpha$ is a ranking function.*

Proof. Note that the normalizing factor δ will ensure that $\kappa \diamond \alpha$ is a ranking function as long as there exists a world $w \in W$ for which $\kappa^{\diamond}_{\alpha}(w) \neq \infty$. It must thus be shown that if there exists an event $e^r \in Evt$ which is event-rational and α is an e^r-signal, then there must exist a world $w \in W$ for which $\kappa^{\diamond}_{\alpha}(w) \neq \infty$.

κ is assumed to be a ranking function. Let w^- be a world for which $\kappa(w^-) \neq \infty$. By definition of the transition function, there must exist a world w^+ for which $T(w^-, e, w^+) \neq \infty$, for all $e \in Evt$. Choose the e^r which is event-rational. Then $E(e^r, w^-) \neq \infty$. Furthermore, because w^- exists such that $T(w^-, e^r, w^+) \neq \infty$ and we know that α is an e^r-signal, $O_Q(\alpha, w^+) \neq \infty$.

By definition of operation $\diamond$ (3), $\kappa^{\diamond}_{\alpha}(w^+) = O_Q(\alpha, w^+) + \min_{\substack{w \in W \\ e \in Evt}} \big\{T_Q(w, e, w^+) + E_Q(e, w) + \kappa(w)\big\} - \delta = O_Q(\alpha, w^+) + \min\big\{\ldots, T_Q(w^-, e^r, w^+) + E_Q(e^r, w^-) + \kappa(w^-), \ldots\big\} - \delta \neq \infty$. □

Proposition 1. *If the HQBC model is specified such that α is strongly observable, there exists an event-rational event $e \in Evt$ possible in κ, and α is an e-signal, then (CM) holds.*

Proof. $\kappa \lhd \alpha := \{(w, n) \mid w \in W, n = \min\{\kappa^{CPI}_{\alpha}(w) + es(\alpha, w), \quad \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\} - \delta''\}$. Recall that neither $es(\alpha, w)$ nor $os(\alpha, w)$ can have a value of ∞. And given the antecedents of the proposition, by Lemmata 1 and 2, there must be a w^0 for which $\kappa^{CPI}_{\alpha}(w^0) = 0$ or $\kappa^{\diamond}_{\alpha}(w^0) = 0$. Hence, either $\kappa^{CPI}_{\alpha}(w^0) + es(\alpha, w^0) \neq \infty$ or $\kappa^{\diamond}_{\alpha}(w^0) + os(\alpha, w^0) \neq \infty$. Thus $\kappa^{\lhd}_{\alpha}(w^0) \neq \infty$ and due to the normalizing factor δ'', there exists a w s.t. $\kappa^{\lhd}_{\alpha}(w) = 0$. □

Postulate (S) If κ is a ranking function, then $\kappa^{\lhd}_{\alpha}(\alpha) = 0$.

Lemma 3. *If α is strongly observable, then $\kappa^{CPI}_{\alpha}(\alpha) = 0$.*

Proof. Omitted to save space; available on request. □

Lemma 4. *Let the HQBC model be specified such that there exists an event-rational event $e \in Evt$ possible in κ, and α is a trustworthy e-signal. Then $\kappa^{\diamond}_{\alpha}(\alpha) = 0$.*

Proof. Recall that $\kappa^{\diamond}_{\alpha}(\alpha) = \min_{w \in W, w \Vdash \alpha} \kappa^{\diamond}_{\alpha}(w)$. By Lemma 2, $\kappa^{\diamond}_{\alpha}$ is a ranking function and thus there exists a w for which $\kappa^{\diamond}_{\alpha}(w) = 0$. Hence, for $\kappa^{\diamond}_{\alpha}(\alpha)$ not to equal 0, there must exist a $w' \in W$ s.t. $w' \not\Vdash \alpha$ and $\kappa^{\diamond}_{\alpha}(w') \neq \infty$. But then $O_Q(\alpha, w') \neq \infty$. Therefore, for (S) not to hold, an agent needs to believe that $O_Q(\alpha, w') \neq \infty$ for some world w' where $w' \not\Vdash \alpha$. But then α cannot be trustworthy. Arguing by contradiction, (S) must hold. □

Note that trustworthiness is required for Lemma 4, in addition to the antecedents required for Lemma 2.

Proposition 2. *If the HQBC model is specified such that α is strongly observable, there exists an event-rational event $e \in Evt$ possible in κ, and α is a trustworthy e-signal, then (S) holds.*

Proof. Note that neither $es(\cdot)$ nor $os(\cdot)$ can have a value of ∞. Moreover, because α is trustworthy, by the definitions of CPI and $\diamond$, $\kappa^{CPI}_{\alpha}(w) = \kappa^{\diamond}_{\alpha}(w) = \infty$ whenever $w \not\Vdash \alpha$. Together with Lemmata 3 and 4, one can thus infer that

$$\begin{aligned}\delta'' &= \min_{w \in W} \{\min\{\kappa^{CPI}_{\alpha}(w) + es(\alpha, w), \quad \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\}\} \\ &= \min_{\substack{w \in W \\ w \Vdash \alpha}} \{\min\{\kappa^{CPI}_{\alpha}(w) + es(\alpha, w), \quad \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\}\}\end{aligned}$$

Then

$$\begin{aligned}\kappa^{\lhd}_{\alpha}(\alpha) &= \min_{\substack{w \in W \\ w \Vdash \alpha}} \{\kappa^{\lhd}_{\alpha}(w)\} \\ &= \min_{\substack{w \in W \\ w \Vdash \alpha}} \{\min\{\kappa^{CPI}_{\alpha}(w) + es(\alpha, w), \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\} - \delta''\} \\ &= \min_{\substack{w \in W \\ w \Vdash \alpha}} \{\min\{\kappa^{CPI}_{\alpha}(w) + es(\alpha, w), \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\}\} - \delta'' \\ &= 0.\end{aligned}$$

□

5 Comparison of HQBC with Generalized Update

Boutilier [3] adopts an event-based approach where a set of events E is assumed. These events are allowed to be nondeterministic, and each possible outcome of an event is ranked according to its plausibility via a ranking function. "As in the original event-based semantics, we will assume each world has an event ordering associated with it that describes the plausibility of various event occurrences at that world," [3, p. 14].

A *generalized update model* is then defined as $\langle W, \kappa, E, \mu \rangle$, where

- W is a set of possible worlds;
- κ is a ranking over W (the agent's epistemic state);
- E is a mapping from $w \in W$ and $e \in Evt$ to rankings $\kappa_{w,e}$ over W, where $\kappa_{w,e}(w')$ describes the plausibility that world w' results when event e occurs at world w;
- μ is a mapping from $w \in W$ to rankings κ_w over Evt, where $\kappa_w(e)$ captures the plausibility of the occurrence of event e at world w.

In this model, the set of events Evt is implicit and the (initial) epistemic state κ explicit.

Lemma 5. *T_Q corresponds to E, and E_Q corresponds to μ. The correspondence is in the sense that the values of functions are equal for the same arguments, respectively, parameters.*

Proof. Omitted to save space; available on request. □

In the rest of the paper, due to Lemma 5, we shall assume that $T_Q(w, e, w')$ and $\kappa_{w,e}(w')$ are interchangeable, and that $E_Q(e, w)$ and $\kappa_w(e)$ are interchangeable.

Boutilier calls the evolution of w into w', under event e, a *transition*, which he writes $w \xrightarrow{e} w'$. He defines (rhs in our notation)

$$\kappa(w \xrightarrow{e} w') := T_Q(w, e, w') + E_Q(e, w) + \kappa(w).$$

And he defines the set of possible α-transitions:

$$\begin{aligned}Tr(\alpha, \kappa) := \{w \xrightarrow{e} w' \mid w, w' \in W, e \in Evt, \\ w' \Vdash \alpha, \kappa(w \xrightarrow{e} w') \neq \infty\}.\end{aligned}$$

$Tr(\alpha, \kappa)$ is the set of transitions from one world to the next via an event, such that the transition (T_Q) is possible, the event in the departure world (E_Q) is possible, and the departure world (κ) is possible, and such that the arrival world is an α-world.

Then Boutilier defines

$$result^{GU}(\alpha, \kappa) := \{w \mid w' \xrightarrow{e} w \in \min Tr(\alpha, \kappa)\}$$

and defines *generalized update* as

$$Bel^{GU}_{\alpha}(\kappa) := \{\beta \in L \mid result^{GU}(\alpha, \kappa) \subseteq [\![\beta]\!]\}. \tag{5}$$

Proposition 3. *A generalized update model can be realized via an HQBC model.*

Proof. Let $G = \langle W, \kappa, E, \mu \rangle$ be a generalized update model, where κ is a (current) epistemic state. Choose an HQBC model $H = \langle W, Evt, T_Q, E_Q, O_Q, os \rangle$ with implicit epistemic state κ and such that, for all $w, w' \in W$ and $e \in Evt$, $T_Q(w, e, w') = \kappa_{w,e}(w')$ and $E_Q(e, w) = \kappa_w(e)$. Then T_Q corresponds to E and E_Q corresponds to μ. G is thus realized via H. □

Lemma 6. *Let $\mathcal{H}$ be the class of HQBC models specified such that there exists an event-rational event $e \in Evt$ possible in κ and such that evidence α is an e-signal, trustworthy and clear. For every generalized update model realizable via an HQBC model in $\mathcal{H}$, $Bel^{GU}_{\alpha}(\kappa) = Bel(\kappa^{\diamond}_{\alpha})$.*

Proof. Let $H \in \mathcal{H}$ s.t. $H = \langle W, Evt, T_Q, E_Q, O_Q, os \rangle$ with implicit epistemic state κ. Let $G = \langle W, \kappa, E, \mu \rangle$ be a generalized update model realized via H. Note that $\kappa(w \xrightarrow{e} w') \neq \infty$ iff $T_Q(w, e, w') \neq \infty$, $E_Q(e, w) \neq \infty$ and $\kappa(w) \neq \infty$.

$Bel^{GU}_{\alpha}(\kappa) = Bel(\kappa^{\diamond}_{\alpha})$ iff $\{\beta \in L \mid result^{GU}(\alpha, \kappa) \subseteq [\![\beta]\!]\} = \{\beta \in L \mid (\kappa^{\diamond}_{\alpha})^{-1}(0) \subseteq [\![\beta]\!]\}$ (by their definitions: (5) and (2)) iff $result^{GU}(\alpha, \kappa) = (\kappa^{\diamond}_{\alpha})^{-1}(0)$ iff $result^{GU}(\alpha, \kappa) = [\![Bel(\kappa^{\diamond}_{\alpha})]\!]$.

And $result^{GU}(\alpha, \kappa) =$

$$\{w \mid w' \xrightarrow{e} w \in \min Tr(\alpha, \kappa)\} \text{ (by definition of } result^{GU}\text{)}$$

$$\begin{aligned}=&\{w \mid w' \xrightarrow{e} w \in \min\{w' \xrightarrow{e} w \mid w', w \in W, \\ &e \in Evt, w \Vdash \alpha, \kappa(w' \xrightarrow{e} w) \neq \infty\}\} \text{ (by definition of } Tr\text{)}\end{aligned}$$

$$= \mathop{\arg\min}_{\substack{w, w' \in W,\, e \in Evt,\, w \Vdash \alpha \\ w:\ T_Q(w', e, w) + E_Q(e, w') \\ + \kappa(w') \neq \infty}} \{T_Q(w', e, w) + E_Q(e, w') + \kappa(w')\}$$

(by definition of $\kappa(w' \xrightarrow{e} w)$)

$$\begin{aligned}= \mathop{\arg\min}_{w \in W} \{&O_Q(\alpha, w) \\ &+ \min_{w' \in W, e \in Evt} \{T_Q(w', e, w) + E_Q(e, w') + \kappa(w')\}\}\end{aligned}$$

(by definition of class $\mathcal{H}$, $w \Vdash \alpha$ and $T_Q(w', e, w) + E_Q(e, w') + \kappa(w') \neq \infty$)

$$= \left\{ \arg\min_{w \in W} \{\kappa^{\diamond}_{\alpha}(w)\} \right\} = (\kappa^{\diamond}_{\alpha}(w))^{-1}(0) = [\![Bel(\kappa^{\diamond}_{\alpha})]\!].$$

□

Proposition 4. *Let $\mathcal{H}$ be the class of HQBC models specified such that there exists an event-rational event $e \in Evt$ possible in κ and such that evidence α is an e-signal, trustworthy and clear. For every generalized update model realizable via an HQBC model in $\mathcal{H}$, $Bel^{GU}_{\alpha}(\kappa) = Bel(\kappa^{\triangleleft}_{\alpha})$.*

Proof. Let $os(\alpha, w) = 0$ for all possible α and for all $w \in W$. Let $\tau > \kappa^{\diamond}_{\alpha}(w) - \kappa^{OGI}_{\alpha}(w)$ for all $w \in W$ for which $\kappa^{\diamond}_{\alpha}(w) \neq \infty$ and $\kappa^{OGI}_{\alpha}(w) \neq \infty$. Recall that τ may not equal ∞ and for all $w \in W$, $os(\alpha, w) \neq \infty$.

Then, for all $w \in W$,

$$\begin{aligned}
\kappa^{\triangleleft}_{\alpha}(w) &= \min\{\kappa^{OGI}_{\alpha}(w) + es(\alpha, w), \quad \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\} - \delta'' \\
&= \min\{\kappa^{OGI}_{\alpha}(w) + \tau - os(\alpha, w), \kappa^{\diamond}_{\alpha}(w) + os(\alpha, w)\} - \delta'' \\
&\quad \text{(by Def. 2)} \\
&= \min\{\kappa^{OGI}_{\alpha}(w) + \tau, \quad \kappa^{\diamond}_{\alpha}(w) + 2os(\alpha, w)\} - \delta'' \\
&= \min\{\kappa^{OGI}_{\alpha}(w) + \tau, \quad \kappa^{\diamond}_{\alpha}(w)\} - \delta'' \\
&\quad \text{(by the above definition of } os(\alpha, w)\text{)} \\
&= \kappa^{\diamond}_{\alpha}(w) - \delta'' \text{ (by the above definition of } \tau\text{)} \\
&= \kappa^{\diamond}_{\alpha}(w) \text{ (by definition, } \min_{w' \in W} \kappa^{\diamond}_{\alpha}(w') = 0\text{).}
\end{aligned}$$

Then $Bel(\kappa^{\triangleleft}_{\alpha}) = Bel(\kappa^{\diamond}_{\alpha})$ and by Lemma 6, $Bel(\kappa^{\diamond}_{\alpha}) = Bel^{GU}_{\alpha}(\kappa)$. □

6 Examples

We use Boutilier's two examples [3, § 3.3]. One can then compare his *generalized update* (GU) with our HQBC.

The first example involves a book (B) which might be inside the house or on the patio. There are three events: it rains, in which case the grass (G) and the patio get wet, the sprinkler comes on, in which case only the grass gets wet, or nothing happens. In this example, events are deterministic. If the book is on the patio, it will get wet when it rains, else not. If the book is inside and the book is dry, it will never get wet. Figure 1 illustrates the prior epistemic state of an agent who believes its book is on the patio and that both the grass and the book are dry ($\kappa(Patio(B) \land Dry(B) \land Dry(G)) = 0$), but if the book is not on the patio, the agent believes it has left it inside ($\kappa(Inside(B) \land Dry(B) \land Dry(G)) = 1$). The other less plausible worlds are omitted. Event plausibility is ranked as $E_Q(null, w) = 0$, $E_Q(rain, w) = 1$, $E_Q(sprinkler, w) = 2$, for all w (a 'global' ordering suitable for all worlds is assumed). This is the only information required for GU.

For HQBC, the observation function (O_Q), ontic strength (os) and its top value (τ), and distance measure (d) are required, in addition. We let all observations be trustworthy and clear. For now, let the agent have no opinion as to whether observations are ontic or epistemic, that is, for all possible α and for all $w \in W$, $os(\alpha, w) = es(\alpha, w) = 1$.[9] d will be defined in accordance with Hamming distance, as before.

[9] This implies that $\tau = 2$.

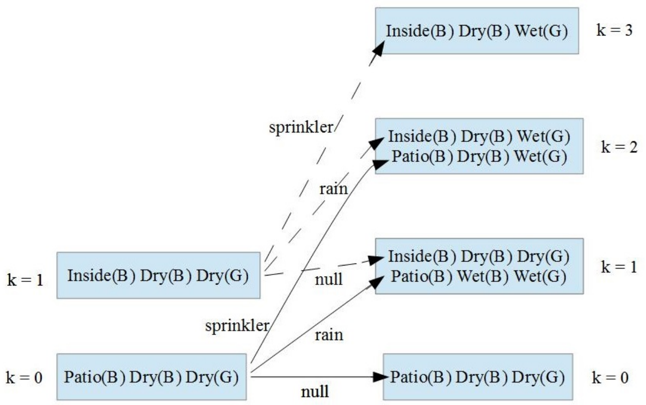

Figure 1: Scenario with multiple events (with deterministic outcomes), including event plausibility information.

Suppose the agent observes that the grass is wet ($\neg Dry(G)$). This contradicts what the agent presently believes, so, with respect to revision, κ CPI $\neg Dry(G)$ is interpreted as κ OGI $\neg Dry(G)$. We determine that

- $\kappa^{\mathsf{OGI}}_{\neg Dry(G)}(Patio(B) \land Dry(B) \land \neg Dry(G)) = 0$
- $\kappa^{\mathsf{OGI}}_{\neg Dry(G)}(\neg Patio(B) \land Dry(B) \land \neg Dry(G)) = 1$
- $\kappa^{\mathsf{OGI}}_{\neg Dry(G)}(w) = \infty$ for all $w \in W$ s.t. $w \nVdash Patio(B) \land Dry(B) \land \neg Dry(G)$ and $w \nVdash \neg Patio(B) \land Dry(B) \land \neg Dry(G)$

With respect to update, we determine that

- $\kappa^{\diamond}_{\neg Dry(G)}(Patio(B) \land Dry(B) \land \neg Dry(G)) = 1$
- $\kappa^{\diamond}_{\neg Dry(G)}(Patio(B) \land \neg Dry(B) \land \neg Dry(G)) = 0$
- $\kappa^{\diamond}_{\neg Dry(G)}(w) = \infty$ for all $w \in W$ s.t. $w \nVdash Patio(B) \land Dry(B) \land \neg Dry(G)$ and $w \nVdash Patio(B) \land \neg Dry(B) \land \neg Dry(G)$

Then combining these results gives

- $\kappa^{\triangleleft}_{\neg Dry(G)}(Patio(B) \land Dry(B) \land \neg Dry(G)) = 0$
- $\kappa^{\triangleleft}_{\neg Dry(G)}(Patio(B) \land \neg Dry(B) \land \neg Dry(G)) = 0$
- $\kappa^{\triangleleft}_{\neg Dry(G)}(\neg Patio(B) \land Dry(B) \land \neg Dry(G)) = 1$

and the other worlds are deemed impossible. If the agent were to reflect on its new beliefs, it might reason as follows.

> I believe $Patio(B) \land Dry(B) \land \neg Dry(G)$ because it is the $\neg Dry(G)$-world closest to my prior beliefs (and at least plausible, because it is plausibly explained by the sprinkler coming on). I believe $Patio(B) \land \neg Dry(B) \land \neg Dry(G)$ because it is a $\neg Dry(G)$-world best explained by rain (in my prior beliefs). I don't fully believe $\neg Patio(B) \land Dry(B) \land \neg Dry(G)$, but it is plausible because it is the $\neg Dry(G)$-world second closest to my prior beliefs (although $\neg Patio(B)$ was previously not fully believed, it was deemed plausible.).

Now suppose the ontic strength of $\neg Dry(G)$ is defined as $os(\neg Dry(G), w) = 0$, for all $w \in W$, and $\tau = 2$. That is, perceiving wet grass is always deemed slightly more ontic than epistemic. Then the resulting epistemic state is determined as in Table 1. In the table, worlds are identified by three letters, such that, for instance, $pdd \models Patio(B) \land Dry(B) \land Dry(G)$ and $iww \models \neg Patio(B) \land \neg Dry(B) \land \neg Dry(G)$; OGI abbreviates $\kappa^{\mathsf{OGI}}_{\alpha}(w)$, $\diamond$ abbreviates $\kappa^{\diamond}_{\alpha}(w)$, es abbreviates $es(\alpha, w)$ and os abbreviated

$os(\alpha, w)$, where α is the incumbent observation and w is the world of the row. The "min" column indicates the minimum value between the two columns to its left, and is actually the rank assigned to the world w of the incumbent row ($\kappa_\alpha^{\triangleleft}(w)$).

Table 1: Agent prefers an ontic interpretation.

World	OGI $+ es$	$\diamond + os$	min
pdd	$\infty + 2$	$\infty + 0$	∞
pdw	$0 + 2$	$1 + 0$	1
pwd	$\infty + 2$	$\infty + 0$	∞
pww	$\infty + 2$	$0 + 0$	0
idd	$\infty + 2$	$\infty + 0$	∞
idw	$1 + 2$	$\infty + 0$	3
iwd	$\infty + 2$	$\infty + 0$	∞
iww	$\infty + 2$	$\infty + 0$	∞

Finally, suppose ontic strength of $\neg Dry(G)$ is defined as $os(\neg Dry(G), w) = 2$ for all $w \in W$, with $\tau = 2$ (which implies that $es(\neg Dry(G), w) = 0$, for all $w \in W$). That is, perceiving wet grass is always deemed more epistemic than ontic. Then the resulting epistemic state is determined as in Table 2.

Table 2: Agent prefers an epistemic interpretation.

World	OGI $+ es$	$\diamond + os$	min
pdd	$\infty + 0$	$\infty + 2$	∞
pdw	$0 + 0$	$1 + 2$	0
pwd	$\infty + 0$	$\infty + 2$	∞
pww	$\infty + 0$	$0 + 2$	2
idd	$\infty + 0$	$\infty + 2$	∞
idw	$1 + 0$	$\infty + 2$	1
iwd	$\infty + 0$	$\infty + 2$	∞
iww	$\infty + 0$	$\infty + 2$	∞

We now analyze the results of the two tables/cases a little.

We see that when the agent prefers to interpret or explain $\neg Dry(G)$ as an ontic observation, the agent considers world pww as most plausible, that is, it fully believes that the book is on the patio, the book is wet and the grass is wet. A reason could be that, given the agent's most plausible prior belief that the book is on the patio and dry and the grass is dry, it rained. This is the same result produced by generalized update [3, p. 17]; this correspondence makes sense, given that our update ($\diamond$) is 'aligned' with GU ($Bel_\alpha^{GU}(\kappa) = Bel(\kappa_\alpha^\diamond)$ under reasonable conditions; Lem. 6). Notice that the plausibility of pww due to revision is not in contention, because $\kappa_{\neg Dry(G)}^{\mathsf{OGI}}(pww) = \infty$.

We see that when the agent prefers to interpret or explain $\neg Dry(G)$ as an epistemic observation, the agent considers world pdw as most plausible, that is, it fully believes that the book is on the patio, the book is dry and the grass is wet. It can be seen from Table 2 that it is revision which causes the agent to believe pdw. Notice that pdw is the Hamming-closest $\neg Dry(G)$-world to the most plausible prior belief (pdd).

The second example is shown in Figure 2. Here only one possible event is assumed, the action of dipping litmus paper in a beaker.

> The beaker is believed to contain either an acid or a base ($\kappa = 0$); little plausibility ($\kappa = r$) is accorded the possibility that it contains some other substance (say, kryptonite). The expected outcome of the test is a color change of the litmus paper: it changes from yellow to red if the substance is an acid, to blue if it is a base, and to green if it is kryptonite. However, the litmus test can fail some small percentage of the time, in which case the paper also turns green. This outcome is also accorded little plausibility ($\kappa = g$). If the paper is dipped, and red is observed, the agent will adopt the new belief $acid$. Unlike KM update [of Katsuno and Mendelzon [8]], generalized update permits observations to rule out possible transitions, or previously epistemically possible worlds. As such, it is an appropriate model for revision and expansion of beliefs due to information-gathering actions. An observed outcome of $green$ presents two competing explanations: either the test failed (the substance is an acid or a base, and we still dont know which) or the beaker contains kryptonite. The most plausible explanation and the updated epistemic state depend on the relative magnitudes of g and r. The figure suggests that $g < r$, so the a test failure is most plausible and the belief $acid \vee base$ is retained. If test failures are more rare ($r < g$), then this outcome would cause the agent to believe the beaker held kryptonite. [3, p. 18]

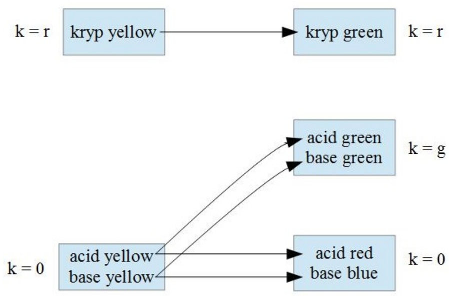

Figure 2: Scenario with single event (with non-deterministic outcomes), including event plausibility information.

Now we investigate how HQBC deals with this scenario for two observations. We let all observations be trustworthy and clear, and Hamming distance is used to define d. The three possible observations are red, $blue$ and $green$.

Tables 3 and 4 show the agent's new epistemic state after perceiving the litmus paper turning red, respectively, green. In both tables, the three right-most columns report the new state ($\kappa_\alpha^{\triangleleft}$) when the agent (from left to right) (i) is indifferent about whether the observation is ontic or epistemic, (ii) prefers an ontic interpretation, (iii) prefers an epistemic interpretation. "$os = x$" ("$es = x$") in a column heading means that $os(\alpha, w) = x$ (resp., $es(\alpha, w) = x$) for all $w \in W$. In the tables, worlds are identified by two letters, such that, for instance, $ar \models acid \wedge red$, $ab \models acid \wedge blue$, $bg \models base \wedge green$, $ky \models krypt \wedge yellow$. To save space, rows containing ∞ in every row are omitted. Of course, perceiving red, blue or green is inconsistent with the current belief that the litmus paper is yellow; revision operator CPI is thus interpreted as OGI.

Table 3: Agent perceives the litmus paper turning red.

World	OGI	$\diamond$	$es = 1$ $os = 1$	$es = 2$ $os = 0$	$es = 0$ $os = 2$
ar	0	0	0	0	0
br	0	∞	0	2	0
kr	r	∞	r	$r + 2$	r

We see that when the agent prefers to revise its beliefs ($es = 0$, $os = 2$), and when it is indifferent about whether to revise or update ($es = 1$, $os = 1$), then its resulting beliefs seem unintuitive to us humans—the agent believes as equally plausible that the substance is acid and that it is base. However, when the agent prefers to update its beliefs ($es = 2$, $os = 0$), then it reasonably believes (only) that the substance is acid. A reasonable agent *should* prefers to update its beliefs because it should consider all its observations in this scenario to be ontic, due to the ontic nature of dipping litmus paper.

Table 4: Agent perceives the litmus paper turning green.

World	OGI	$\diamond$	$es = 1$, $os = 1$	$es = 2$, $os = 0$	$es = 0$, $os = 2$
ag	0	g	0	0	0
bg	0	g	0	0	0
kg	r	$r+r$	r	r	r

Note that for the case when $es = 2$ and $os = 0$, the values had to be normalized ($\delta'' = 2$). It is interesting to see that no matter what stance the agent takes on ontic/epistemic strength, when it perceives $green$, it believes with equal plausibility that the substance is acid and base. Assuming $r > 0$, the substance is less plausibly kryptonite, but not impossible.

In the cases when the agent has an event-based attitude (i.e., it prefers to interpret observations ontically), all results when HQBC is applied align with the results when GU is applied (w.r.t. the examples in this paper).

7 Concluding Remarks, Related and Future Work

A hybrid qualitative belief change (HQBC) construction was presented—based on ranking theory and which trades revision off with update, according to the agent's confidence for whether the received observation/evidence is ontic (due to a physical event) or epistemic (due to an announcement). We proved that HQBC is, in a particular sense, an extension of Boutilier's generalized update (GU). In other words, the HQBC model in a class of 'reasonable' models can be specified to perform exactly the same belief change as GU would. Moreover, HQBC allows for more sophisticated belief change than GU, in particular, with respect to rankic belief revision (based on conditional plausibility and generalized imaging, for instance) and with respect to employing a notion of ontic strength. The examples in this paper support our propositions concerning the relationship between HQBC and GU.

> Determining $os(\alpha, w)$ for every foreseen α in every possible world w will be challenging for a designer. Some deep questions are: Should the designer/agent provide the strengths (via stored values or programmed reasoning), or do these strengths come to the agent attached to the new information? What is the reasoning process we go through to determine whether information is epistemic or ontic, if at all? In general, how does an agent know when information is epistemic (requiring revision) or ontic (requiring update)? [11]

One direction to investigate as a possible answer to the questions above is to condition ontic/epistemic strength on particular propositions. For instance, the more plausible the proposition, the more likely that the received information is ontic. For such an approach to work, the framework would presumably have to accommodate the specification of *condition propositions* for every observation of interest. Revision and update would then be traded off depending on the plausibility/probability of the condition of the observation under consideration.

Friedman and Halpern [5] investigate belief revision and update employing a framework based on time-stamps and runs of possible evolutions of a system. They provide some interesting insights regarding the relationship between revision and update, which may also benefit our future work. In their concluding section, they hint that their framework could 'mix' revision and update: "In this framework, belief change operations can be determined by choosing a plausibility measure that captures the agent's preferences among sequences of worlds.... [T]here are prior plausibilities that, when conditioned on a surprising observation, allow the agent to revise some earlier beliefs and to assume that some change has occurred", [5]. To our knowledge, they never did investigate the 'mixture' approach.

Relationships to the change operations defined by Beierle and Kern-Isberner [1, 2], which make use of knowledge bases, also need to be investigated.

Although Lang's work [**?**] is not directly applicable to ours in terms of 'mixing' revision and update, it does unpack and highlight several important characteristics of update. Lang also discusses the relationship between update to revision. His insights might well guide our future efforts in this area. He writes

> In complex environments, especially planning under incomplete knowledge, actions are complex and have both ontic and epistemic effects; the belief change process then is very much like the feedback loop in partially observable planning and control theory: perform an action, project its effects on the current epistemic state, then get the feedback, and revise the projected epistemic state by the feedback. Clearly, update allows for projection only. Or, equivalently, if one chooses to separate the ontic and the epistemic effects of actions, by having two disjoint sets of actions (ontic and epistemic), then ontic actions lead to projection only, while epistemic actions lead to revision only. Therefore, if one wants to extend belief update so as to handle feedback, there is no choice but integrating some kind of revision process, as in several recent works [...] [**?**]

This act-update-perceive-revise "feedback loop" is the default approach when complex actions/events are considered; it is fundamentally different to the simultaneous, hybrid belief change approach. Yet, we have not come across a convincing argument against the hybrid approach. It seems that the traditional "feedback loop" approach assumes that there is always certainty about the ontic/epistemic status of every piece of information received. A major question for future research is, Is there a theory or framework to synthesize the two approaches?

Nayak [10] proves that, given an appropriate function for measuring distance between worlds, classical revision ($*$) can be reduced to classical update ($\diamond$). Formally, he proves that $(x * k) \diamond x = k * x$, where $k, x \in L$, k is an agent's knowledge and x is the (new) evidence. Nayak points out that the "nice storyline that cleanly demarcates revision from update appears not to be such a good story after all," [10]. Nayak's surprising result is just one more reason to investigate the hybrid belief change approaches.

"We can regard imaging as a probabilistic version of update, and conditionalization as a probabilistic version of revision," [8]. And according to Nayak, KM-update "is known to be the non-probabilistic counterpart of the account of [probabilistic] imaging propounded by David Lewis in order to develop a theory of conditionals [...]" [10]. Dubois and Prade [4] give a version of imaging for belief update in the possibilistic framework. Rens [11] uses imaging (GI) on the revision side; his justification is because imaging can deal with contradictory evidence, whereas conditioning cannot. We are not convinced that imaging is strictly an update process. Where exactly imaging lies on the revision-update spectrum is, to our minds, another deep question still to be answered.

REFERENCES

[1] C. Beierle and G. Kern-Isberner, 'On the modelling of an agent's epistemic state and its dynamic changes', *Electronic Communications of the European Association of Software Science and Technology*, **12**, (2008).

[2] C. Beierle and G. Kern-Isberner, 'Towards an agent model for belief management', in *Advances in Multiagent Systems, Robotics and Cybernetics: Theory and Practice. (Volume III)*, eds., G. Lasker and J. Pfalzgraf, IIAS, Tecumseh, Canada, (2009).

[3] C. Boutilier, 'A unified model of qualitative belief change: A dynamical systems perspective', *Artif. Intell.*, **98**(1–2), 281–316, (1998).

[4] D. Dubois and H. Prade, 'Belief revision and updates in numerical formalisms: An overview, with new results for the possibilistic framework', in *Proceedings of the Thirteenth Intl. Joint Conf. on Artif. Intell. (IJCAI-93)*, pp. 620–625, San Francisco, CA, USA, (1993). Morgan Kaufmann Publishers Inc.

[5] N. Friedman and J. Halpern, 'Modeling belief in dynamic systems. Part II: Revision and update', *Journal of Artif. Intell. Research (JAIR)*, **10**, 117–167, (1999).

[6] P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, Massachusetts/England, 1988.

[7] M. Goldszmidt and J. Pearl, 'Qualitative probabilities for default reasoning, belief revision, and causal modeling', *Artificial Intelligence*, **84**, 57–112, (1996).

[8] H. Katsuno and A. O. Mendelzon, 'On the difference between updating a knowledge base and revising it', in *Belief Revision*, ed., P. Gärdenfors, 183–203, Cambridge University Press, (1992).

[9] D. Lewis, 'Probabilities of conditionals and conditional probabilities', *Philosophical Review*, **85**(3), 297–315, (1976).

[10] A. C. Nayak, 'Is revision a special kind of update?', in *AI 2011: Advances in Artif. Intell.: Proceedings of the Twenty-fourth Australasian Joint Conf.*, LNAI, pp. 432–441, Berlin/Heidelberg, (2011). Springer-Verlag.

[11] G. Rens, 'On stochastic belief revision and update and their combination', in *Proceedings of the Sixteenth Intl. Workshop on Non-Monotonic Reasoning (NMR)*, eds., G. Kern-Isberner and R. Wassermann, pp. 123–132. Technical University of Dortmund, (2016).

[12] G. Rens, T. Meyer, and G. Casini, 'Revising incompletely specified convex probabilistic belief bases', in *Proceedings of the Sixteenth Intl. Workshop on Non-Monotonic Reasoning (NMR)*, eds., G. Kern-Isberner and R. Wassermann, pp. 133–142. Technical University of Dortmund, (2016).

[13] W. Spohn, 'Ordinal conditional functions: A dynamic theory of epistemic states', in *Causation in Decision, Belief Change, and Statistics*, eds., W. Harper and B. Skyrms, volume 42 of *The University of Western Ontario Series in Philosophy of Science*, 105–134, Springer Netherlands, (1988).

[14] W. Spohn, 'A survey of ranking theory', in *Degrees of Belief*, eds., F. Huber and C. Schmidt-Petri, 185–228, Springer Netherlands, Dordrecht, (2009).